

WHAT IS CLAIMED IS:

1. A method for forwarding packet-based traffic through a network node, comprising:
dedicating a group of queues in a network node to a customer;
5 performing queue-specific rate shaping on the customer's traffic according to queue-specific bandwidth limitations respectively associated with the queues; and
performing group-specific rate shaping on the customer's traffic according to a group-specific bandwidth limitation associated with the group of queues.
- 10 2. The method of claim 1, further comprising associating queues from said group of queues with different types of traffic that are to be received from the customer.
3. The method of claim 1, further comprising:
receiving a traffic type bandwidth limitation from the customer; and
15 translating the traffic type bandwidth limitation to a queue-specific bandwidth limitation of said respective queue-specific bandwidth limitations.
4. The method of claim 1, further comprising associating said group of queues with a group rate shaper that performs said group-specific rate shaping on the customer's traffic
20 on an aggregate basis.
5. The method of claim 1, further comprising prioritizing the queues of said group of queues.
- 25 6. The method of claim 5, further comprising:
distributing said portion of excess unused bandwidth among the group of queues on a priority basis according to said prioritizing.

7. The method of claim 1, further comprising:

scheduling packets for forwarding from one or more of said queues in said group of queues, wherein bandwidth consumed by the packets from each of the queues is less than or equal to respective queue-specific bandwidth limitations for the queues;

5 identifying excess unused bandwidth when the consumed bandwidth is less than said group-specific bandwidth limitation; and

distributing a portion of the excess unused bandwidth to a first queue of the group of queues, wherein the sum of the consumed bandwidth and the portion of the excess unused bandwidth is less than or equal to a group-specific bandwidth limitation for the
10 group.

8. A network node for forwarding packet-based traffic, comprising:

a plurality of queues;

a plurality of queue-specific rate shapers respectively associated with the plurality of

15 queues;

a plurality of group-specific rate shapers configured to be associated with groups of the plurality of queues; and

a group establishment module configured to dedicate a group of said queues to a customer and to associate one of said group-specific rate shapers with said group of
20 queues that is dedicated to said customer.

9. The device of claim 8, further comprising:

a scheduler, coupled to the plurality of queue-specific rate shapers and the plurality of group-specific rate shapers, configured to schedule packets enqueued in the plurality of
25 queues according to the respective plurality of queue-specific rate shapers, wherein the queue-specific rate shaper respectively associated with each queue is associated with a priority, and wherein the scheduler schedules according to the associated priority.

10. The device of claim 9, wherein said scheduler is further configured to:

30 scheduling packets for forwarding from a first one or more queues of said plurality of queues, wherein bandwidth consumed by the packets from each of the first one or more

queues is less than or equal to respective queue-specific bandwidth limitations for the first one or more queues;

identifying excess unused bandwidth when the consumed bandwidth is less than a group-specific bandwidth limitation, wherein a sum of the consumed bandwidth and the excess unused bandwidth approximately equals the group-specific bandwidth limitation;
5 and

scheduling packets for forwarding from a second one or more queues of said plurality of queues using the excess unused bandwidth.

10 11. The device of claim 11, further comprising:

a scheduler configured to:

schedule, in a first round, packets enqueued in the plurality of queues according to the respective plurality of queue-specific rate shapers; and

schedule, in a second round, packets enqueued in the plurality of queues
15 according to the respective group-specific rate shapers.

12. The device of claim 12, wherein said scheduler is further configured to:

schedule, in subrounds of the first round, packets enqueued in the plurality of queues according to a priority respectively associated with each of the queues.

20

13. The device of claim 8, wherein said scheduler is further configured to:

schedule, in subrounds of the second round, packets enqueued in the plurality of queues according to a priority respectively associated with each of the queues.

25 14. The device of claim 8, further comprising:

a plurality of pipes, wherein each pipe is associated with a group-specific rate shaper, and wherein each pipe of said plurality of pipes includes:

multiple traffic channels comprising one or more queues of the plurality of queues, wherein each traffic channel is associated with a queue-specific rate shaper.

30

15. A method for packet-based traffic forwarding, comprising:

dedicating multiple traffic channels to a customer;

performing traffic-type-specific rate shaping according to traffic-type-specific
bandwidth limitations respectively associated with the traffic channels; and

5 performing customer-specific rate shaping according to a customer-specific
bandwidth limitation associated with the traffic channels.

16. The method of claim 15, further comprising:

prioritizing the traffic channels relative to one another.

10

17. The method of claim 16, wherein said performing traffic-type-specific rate shaping
consumes less bandwidth than said customer-specific bandwidth limitation, said method
further comprising:

identifying excess unused bandwidth following the traffic-type-specific rate shaping;

15 and

distributing the excess unused bandwidth to a subset of the traffic channels in priority
order according to said prioritizing.

18. The method of claim 15, further comprising:

20 associating a traffic type with each traffic channel.

19. The method of claim 18, further comprising:

adjusting the traffic-type-specific rate shaping according to traffic type-specific rate
shaping customer preferences.

25

20. The method of claim 15, further comprising:

associating respective traffic-type-specific bandwidth limitations with each traffic
channel such that a sum of the respective traffic-type-specific bandwidth limitations is
less than or equal to the customer-specific bandwidth limitation.

30